



DARM: Control-Flow Melding for SIMT Thread Divergence Reduction

Charitha Saumya, Kirshanthan Sundararajah, and Milind Kulkarni

School of Electrical and Computer Engineering

Purdue University

West Lafayette, IN, USA

cgusthin@purdue.edu, ksundar@purdue.edu, milind@purdue.edu

Abstract—GPGPUs use the Single-Instruction-Multiple-Thread (SIMT) execution model where a group of threads—wavefront or warp—execute instructions in lockstep. When threads in a group encounter a branching instruction, not all threads in the group take the same path, a phenomenon known as control-flow divergence. The control-flow divergence causes performance degradation because both paths of the branch must be executed one after the other. Prior research has primarily addressed this issue through architectural modifications. We observe that certain GPGPU kernels with control-flow divergence have similar control-flow structures with similar instructions on both sides of a branch. This structure can be exploited to reduce control-flow divergence by melding the two sides of the branch allowing threads to reconverge early, reducing divergence. In this work, we present DARM, a compiler analysis and transformation framework that can meld divergent control-flow structures with similar instruction sequences. We show that DARM can reduce the performance degradation from control-flow divergence.

Index Terms—GPGPUs, Control-Flow Divergence, Compiler Optimizations

I. INTRODUCTION

General Purpose Graphics Processing Units (GPGPU) are capable of executing thousands of threads in parallel, efficiently. Advancements in the programming models and compilers for GPUs have made it much easier to write data-parallel applications. Unfortunately, exploiting data parallelism does not immediately translate to better performance. One key reason for the lack of performance portability is that GPGPUs are not capable of executing all the threads independently. Instead threads are grouped together into units called *warps*, and threads in a warp execute instructions in lockstep. This is commonly referred to as the Single Instruction Multiple Thread (SIMT) execution model.

The SIMT model suffers performance degradation when threads exhibit *irregularity* and can no longer execute in lock-step. Irregularity comes in two forms, irregularity in memory accesses patterns (*i.e.* memory divergence) and irregularity in the control-flow of the program (*i.e.* control-flow divergence). Memory divergence occurs when GPGPU threads need to access memory at non-uniform locations, which results in un-coalesced memory accesses. Un-coalesced memory accesses are bad for GPU performance because memory bandwidth can not be fully utilized to do useful work.

TABLE I
COMPARISON OF TECHNIQUES FOR DIVERGENCE REDUCTION

Control-flow and instruction Pattern	Technique		
	Tail Merging	Branch Fusion	DARM
Diamond control-flow with identical instruction sequences	✓	✓	✓
Diamond control-flow with distinct instruction sequences	✗	✓	✓
Complex control-flow	✗	✗	✓

Control-flow divergence occurs when threads in a warp diverge at branch instructions. At the *diverging* branch, lock-step execution can not be maintained because threads in a warp may want to execute different basic blocks (*i.e.* diverge). Instead, when executing instructions along a diverged path, GPGPUs mask out the threads that do not want to take that path. The threads *reconverge* at the Immediate Post-DOMinator (IPDOM) of a divergent branch—the instruction that all threads from both branches want to execute. This style of IPDOM-based reconvergence is implemented in hardware in most GPGPU architectures to maintain SIMT execution. Even though IPDOM-based reconvergence can handle arbitrary control-flow, it imposes a significant performance penalty if a program has a lot of divergent branches. In the IPDOM reconvergence model, instructions executed on divergent branches necessarily cannot utilize the full width of a SIMD unit. If the code has a lot of nested divergent branches or divergent branches inside loops, this style of execution causes significant under-utilization of SIMD resources.

For some GPGPU applications divergent branches are unavoidable, and there have been many techniques proposed to address this issue both in hardware and software. Proposals such as Dynamic warp formation [1], Thread block compaction [2] and Dual-path execution [3] focus on mitigating the problem at the hardware level by changing how threads are scheduled for execution and making sure that threads following the same path are grouped together. Unfortunately, such approaches are not useful on commodity GPGPUs.

There have also been efforts to reduce divergence through compiler approaches that leverage the observation that different control-flow paths often contain similar instruction (sub)sequences. *Tail merging* [4] identifies branches that have identical sequences of code and introduces early jumps to

merged basic blocks, with the effect of reducing divergence. *Branch fusion* generalizes tail merging to work with instruction sequences that may not be identical [5]. However, branch fusion cannot analyze complex control-flow and hence it is restricted to simple *if-then-else* branches where each path has a single basic block (*i.e.* diamond-shaped control-flow).

This paper introduces a more general, software-only approach of exploiting similarity in divergent paths, called *control-flow melding*. Control-flow melding is a general control-flow transformation which can meld similar control-flow *subgraphs* inside a *if-then-else* region (not just individual basic blocks). By working hierarchically, recursively melding divergent control-flow at the level of subgraphs of the CFG, control-flow melding can handle substantially more general control structures than prior work. This paper describes DARM, a realization of control-flow melding for general GPGPU programs. Table I compares the capabilities of DARM with branch fusion and tail merging.

DARM works in several steps. First, it detects divergent *if-then-else* regions and splits the divergent regions into Single Entry Single Exit (SESE) control-flow subgraphs. Next it uses a hierarchical sequence alignment technique to *meld* profitable control-flow subgraphs, repeatedly finding subgraphs whose control-flow structures and constituent instructions can be aligned. Once a fixpoint is reached, DARM uses this hierarchical alignment to generate code for the region with reduced control-flow divergence.

The main contributions of the paper are,

- Divergence-Aware-Region-Melder (DARM), a realization of control-flow melding that identifies profitable melding opportunities in divergent *if-then-else* regions of the control-flow using a hierarchical sequence alignment approach and then melds these regions to reduce control-flow divergence.
- An implementation of DARM in LLVM [6] that can be applied to GPGPU programs written in HIP [7] or CUDA [8]. Our implementation of DARM is publicly available as an archival repository¹ and up-to-date version is available in GitHub².
- An evaluation of DARM on a set of synthetic GPU programs and a set of real-world GPU applications showing its effectiveness

II. BACKGROUND

A. GPGPU Architecture

Modern GPGPUs have multiple processing cores, each of which contains multiple parallel lanes (*i.e.* SIMD units), a vector register file and a chunk of shared memory. The unit of execution is called a warp (or frontend). A warp is a collection of threads executed in lock-step on a SIMD unit. Shared memory is shared among the warps executing on a core. A branch unit takes care of control-flow divergence by maintaining a SIMT stack to enforce IPDOM based reconvergence, as discussed in Section I. GPGPU programming

```

1  __global__ static void bitonicSort(int *values) {
2      // copy data from global memory to shared memory
3      __syncthreads();
4      for (unsigned int k = 2; k <= NUM; k *= 2) {
5          for (unsigned int j = k / 2; j > 0; j /= 2) {
6              unsigned int ixj = tid ^ j;
7              if (ixj > tid) {
8                  if ((tid & k) == 0) {
9                      if (shared[ixj] < shared[tid])
10                         swap(shared[tid], shared[ixj]);
11                 }
12                 else {
13                     if (shared[ixj] > shared[tid])
14                         swap(shared[tid], shared[ixj]);
15                 }
16             }
17             __syncthreads();
18         }
19     }
20 } // write data back to global memory

```

Fig. 1. Bitonic sort kernel

abstractions like CUDA [8] or HIP [7] gives the illusion of data parallelism with independent threads. However, during real execution, a group of program instances (*i.e.* threads) are mapped to a warp and executed in lock-step. Therefore control-flow divergence in SPMD programs is detrimental to the performance because of the SIMT execution limitations.

B. LLVM SSA Form and GPU Divergence Analysis

LLVM [6] is a general framework for building compilers, optimizations and code generators. Most of the widely adopted GPGPU compilers [9], [10] are built on top of the LLVM infrastructure. LLVM uses a target-independent intermediate representation, LLVM-IR, that enables implementing portable compiler optimizations. LLVM-IR uses static single assignment form [11] which requires that every program variable is assigned once and is defined before being used. SSA form uses ϕ nodes to resolve data-flow when branches are present, selecting which definition should be chosen at a confluence of different paths. In GPGPU compilers, a key step in identifying divergent control-flow regions is performing compiler analyses to identify divergent variables (or branches) [5], [12]. A branch is divergent if the branching condition evaluates to a non-uniform value for different threads in a warp. If the branching condition is divergent, threads in a warp will have to take different control-flow paths at this point. LLVM's divergence analysis tags a branch as divergent, if the branching condition is either data-dependent or sync-dependent on a divergent variable (such as thread ID) [12], though more sophisticated divergence analyses have been proposed [13].

III. MOTIVATING EXAMPLE

Bitonic sort is a kernel used in many parallel sorting algorithms such as bitonic merge sort and Cederman's quicksort [14], [15]. Figure 1 shows a CUDA implementation of bitonic sort. This kernel is our running example for describing DARM's control-flow melding algorithm.

¹<https://doi.org/10.5281/zenodo.5784768>

²<https://github.com/charitha22/cgo22ae-darm-code>

In this kernel, the branch condition at line 8 depends on the *thread ID*. Therefore it is divergent. Since the divergent branch is located inside a loop, the execution of the two sides of the branch needs to be serialized many times, resulting in high control-flow divergence. However the code inside the *if* (line 9-10) and *else* (line 13-14) sections of the divergent branch are similar in two ways. First, both code sections have the same control-flow structure (*i.e.* *if-then* branch). Second, instructions along the two paths are also similar. Both conditions compare two elements in the *shared* array and perform a *swap* operation. Therefore the contents of the *if* and *else* sections can be melded to reduce control-flow divergence. Both code sections consists of shared memory loads and store operations. In the unmelded version of the code these shared memory operations will have to be serialized due to thread-divergence. However, if the two sections are melded threads can issue the memory instructions in the same cycle resulting in improved performance.

Existing compiler optimizations such as tail merging and branch fusion cannot be applied to this case. Tail merging is applicable only if two basic blocks have a common destination and have identical instruction sequences at their tails. However in bitonic sort, the *if* and *then* sections of the divergent branch have multiple basic blocks, and the compiler cannot apply tail merging. Similarly branch fusion requires diamond shaped control-flow and does not work if the *if* and *else* sections of the branch contain complex control-flow structures.

DARM solves this problem in two phases. In the analysis phase (Section IV-C), DARM analyzes the control-flow region dominated by a divergent branch to find isomorphic sub-regions that are in the true and false paths of the divergent branch. These isomorphic sub-region pairs are aligned based on their melding profitability using a sequence alignment strategy. Melding profitability is a compile-time approximation of the percentage of thread cycles that can be saved by melding two control-flow regions. Next, DARM chooses profitable sub-region pairs in the alignment (using a threshold) and computes an instruction alignment for corresponding basic blocks in the two regions. In the code generation phase (Section IV-D), DARM uses this instruction alignment to meld corresponding basic blocks in the sub-region pair. This melding is applied iteratively until no further profitable melding can be performed. DARM’s melding transformation is done in SSA form, therefore the resulting CFG can be optimized further using other compiler optimizations (Sections IV-E and IV-F).

IV. DETAILED DESIGN

In this section we describe the algorithm used by DARM to meld similar control-flow subgraphs. First we define the following terms used in our algorithm description.

A. Preliminaries and Definitions

Definition 1. Simple Region : A simple region is a subgraph of a program’s CFG that is connected to the remaining CFG with only two edges, an entry edge and an exit edge.

Definition 2. Region : A region of the CFG is characterized by two basic blocks, its entry and exit. All the basic blocks inside a region are dominated by its entry and post-dominated by its exit. Region with entry E and exit X is denoted by the tuple (E, X) . LLVM regions are defined similarly [16], [17].

Definition 3. Single Entry Single Exit Subgraph : Single entry single exit (SESE) subgraph is either a simple region or a single basic block with a single predecessor and a successor.

Note that a region with entry E and exit X can be transformed into a simple region by introducing a new entry and exit blocks E_{new} , X_{new} . All successors of E are moved to E_{new} and E_{new} is made the single successor of E . Similarly, all predecessors of X are moved to X_{new} and a single exit edge is added from X_{new} to X .

Definition 4. Simplified Region : A region with all its sub-regions transformed into simple regions is called a simplified region.

We now turn to the steps the DARM compiler pass takes to reduce control divergent code.

B. Detecting Meldable Divergent Regions

First DARM needs to detect divergent branches in the CFG. We use LLVM’s built-in divergence analysis to decide if a branch is divergent or not (Section II). The smallest CFG region enclosing a divergent branch is called the *divergent region* corresponding to this branch. Melding transformation is applied only to divergent regions of the CFG. The next step is to decide if a divergent region contains control-flow subgraphs (definition 3) that can be safely melded.

Definition 5. Meldable Divergent Region: A simplified region R with entry E and exit X is said to be *meldable and divergent* if the following conditions are met,

- 1) The entry block of R has a divergent branch
- 2) Let B_T and B_F be the successor blocks of E . B_T does not post-dominate B_F and B_F does not post-dominate B_T

According to definition 5, a *meldable divergent region* has a divergent branch at its entry (condition 1). This makes sure that our melding transformation is only applied to divergent regions, and non-divergent parts of the control-flow are left untouched. Condition 2 ensures that paths $B_T \rightarrow X$ (*i.e.* true path) and $B_F \rightarrow X$ (*i.e.* false path) consists of at least one SESE subgraph and these subgraphs from the two paths can potentially be melded to reduce control-flow divergence. Consider our running example in Figure 1. When this kernel is compiled with *ROCm HIPCC* GPU compiler [7] with *-O3* optimization level into LLVM-IR, we get the CFG shown in Figure 4a. Note that the compiler aggressively unrolls both the loops (lines 4 and 5) in the kernel, and the resulting CFG consists of multiple repeated segments of the inner loop’s body (lines 6-17). In Figure 4a, only one unrolled instantiation of the loop body is shown. As explained in Section III, this kernel contains a divergent branch, which is at the end of basic block

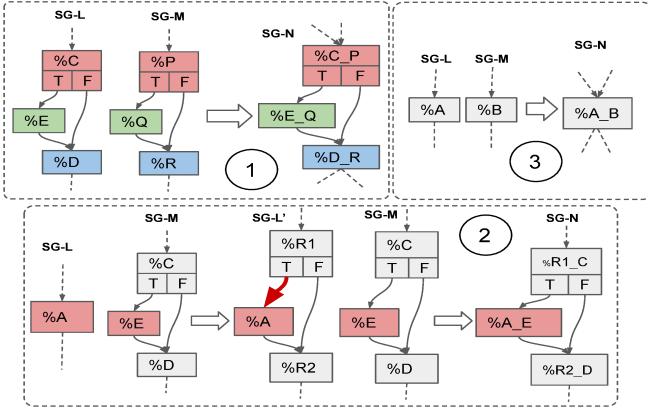


Fig. 2. Examples showing the 3 cases considered by DARM to detect meldable subgraphs

%B. Also %B's two successors %C and %D do not post-dominate each other. Therefore the region (%B, %G) is a meldable divergent region.

C. Computing Melding Profitability

Definition 5 only allows us to detect regions that may contain meldable control-flow subgraphs. It does not tell us whether it is legal to meld them or melding them will improve performance. First we need to define what conditions needs to be satisfied for two SESE subgraphs to be meldable.

Definition 6. Meldable SESE Subgraphs: SESE subgraphs S_1 and S_2 where S_1 belongs to the true path and S_2 belongs to the false path are meldable if any one of the following conditions are satisfied,

- 1) Both S_1 and S_2 have more than one basic block and they are structurally similar i.e. isomorphic.
- 2) S_1 is a simple region and S_2 consists of a single basic block or vice versa.
- 3) Both S_1 and S_2 consists of single basic block.

Definition 6 ensures that any two SESE subgraphs that meets any one of these conditions can be melded without introducing additional divergence to the control-flow. Note that we do not consider subgraphs that contain *warp-level intrinsics* [18] for melding because melding such subgraphs can cause deadlock. Figure 2 shows three examples where each of the above conditions are applicable. Assume in each example subgraphs L and M are in a divergent region (E, X) and only one of the subgraphs are executed from any program path from E to X . (i.e. any thread in warp that executes E must either go through L or M but not both).

Region to Region Melding : In case ①, two SESE subgraphs L and M are isomorphic, therefore they can be melded to have the same control-flow structure (subgraph N in Figure 2-①). In the melded subgraph N , basic blocks %C_P and %D_R are guaranteed to post-dominate E and threads can reconverge at these points resulting in reduction in control-flow divergence. Also the structural similarity in case ① ensures that we do not introduce any additional branches into the melded subgraph.

Basic Block to Region Melding : In case ②, basic block %A (in subgraph L) can potentially be melded with any basic block in CFG M . Assume that basic blocks %A and %E have the most melding profitability (melding profitability described later). First we replicate the control-flow structure of M to create a new CFG L' . Then we place %A in L' such that %A and %E are in similar positions in the the two CFGs L' and M . We also ensure the correctness of the program by concretizing the branch conditions in L' to always execute %A and create ϕ nodes at dominance frontiers of %A to make sure values defined inside %A are reached to their users [11]. In this example branch at end of basic block %R1 will always take the edge %R1 – %A (bold arrow in subgraph L') and ϕ nodes will be added to %R2. Now subgraphs L' and M are isomorphic and therefore can be melded similar to case ①. We refer to this process as *Region Replication*. Main benefit of region replication is that it allows us to meld %A with any profitable basic block in subgraph M and resultant subgraph N has less divergence because threads can reconverge at basic blocks %R1_C and %R2_D in melded subgraph N .

Basic Block to Basic Block Melding : Case ③ is the simplest form where two SESE basic blocks are melded.

A meldable divergent region can potentially have multiple SESE subgraphs in its true and false paths. Therefore we need a strategy to figure out which subgraph pairs to meld. We formulate this as a sequence alignment problem as follows. First, we obtain a ordered sequence of subgraphs in true path and false of the divergent region. Subgraphs are ordered using the post-dominance relation of their entry and exit blocks. For example, if entry node of subgraph S_2 post-dominates exit node of subgraph S_1 , then S_2 comes after S_1 in the order and denoted as $S_1 \prec S_2$. A subgraph alignment is defined as follows,

Definition 7. Subgraph Alignment: Assume a divergent region (E, X) has ordered SESE subgraphs $\{S_1^T, S_2^T, \dots, S_m^T\}$ in its true path and ordered subgraphs $\{S_1^F, S_2^F, \dots, S_n^F\}$ in the false path. A subgraph alignment is an ordered sequence of tuples $A = \{(S_{i0}^T, S_{j0}^F), (S_{i1}^T, S_{j1}^F), \dots, (S_{ik}^T, S_{jk}^F)\}$ where,

- 1) if $(S_p^T, S_q^F) \in A$ then S_p^T and S_q^F are meldable subgraphs
- 2) if $(S_{p1}^T, S_{q1}^F) \prec (S_{p2}^T, S_{q2}^F)$ then $S_{p1}^T \prec S_{p2}^T$ and $S_{q1}^F \prec S_{q2}^F$

According to definition 7, only meldable subgraphs are allowed in a alignment tuple and if the aligned subgraphs are melded, the resultant control-flow graph does not break the original dominance and post-dominance relations of the subgraphs.

Given a suitable alignment scoring function F and gap penalty function W , we can find an optimal subgraph alignment using a sequence alignment method such as Smith-Waterman [19] algorithm. The scoring function F measures the profitability of melding two meldable subgraphs S_1 and S_2 . Prior techniques have employed instruction frequency to approximate the profit of merging two functions [20], [21]. We use a similar method to define subgraph melding profitability. First we define the melding profitability of two basic blocks

$b1$ and $b2$ as follows,

$$MP_B(b1, b2) = \frac{\sum_{i \in Q} \min(freq(i, b1), freq(i, b2)) \times w_i}{lat(b1) + lat(b2)}$$

Here Q is set of all possible instruction types available in the instruction set (*i.e.* LLVM-IR opcodes). $lat(b)$ is the static latency of basic block which can be calculated by summing the latencies of all instructions in b . w_i is the latency of instruction type i . The idea here is to approximate the percentage of instruction cycles that can be saved by melding the instructions in $b1$ and $b2$ assuming a best-case scenario (*i.e.* all common instructions in $b1$ and $b2$ are melded regardless of their order). For example, two basic blocks with identical opcode frequency profile will have a profitability value 0.5.

Because meldable subgraphs are isomorphic, there is a one-to-one mapping between basic blocks (*i.e.* corresponding basic blocks). For example, in Figure 2 case ① the basic block mapping for CFGs L and M are $\{(\%C, \%P), (\%E, \%Q), (\%D, \%R)\}$. Assume the mapping of basic blocks in $S1$ and $S2$ is denoted by O . Subgraph melding profitability MP_S of subgraphs $S1$ and $S2$ is defined in terms of melding profitabilities of their corresponding basic blocks.

$$MP_S(S1, S2) = \frac{\sum_{(b1, b2) \in O} MP_B(b1, b2) \times (lat(b1) + lat(b2))}{\sum_{(b1, b2) \in O} lat(b1) + lat(b2)}$$

Similar to MP_B , MP_S measures the percentage of instruction cycles saved by melding two SESE subgraphs. This metric is an over-approximation, however it provides a fast way of measure the melding profitability of two subgraphs that works well in practice. We use MP_S as the scoring function for subgraph alignment.

Instruction Alignment: Notice that our subgraph melding profitability metric (*i.e.* MP_S) prioritizes subgraph pairs that have many similar instructions in their corresponding basic blocks. Therefore when melding two corresponding basic blocks we must ensure that maximum number of similar instructions are melded together. This requires computing an alignment of two instruction sequences such that if they are melded using this alignment, the number of instruction cycles saved will be maximal. We use the approach used in Branch Fusion [5] to compute an optimal alignment for two instructions sequences. In this approach compatible instructions are aligned together and instructions with higher latency are prioritized to be aligned over lower latency instructions. Compatibility of two instructions for melding depends on a number of conditions like having the same opcode and types of the operands being compatible. We used the criteria described by Rocha et al. [21] to determine this compatibility. This instruction alignment model uses a gap penalty for unaligned instructions because extra branches needs to be generated to conditionally execute these unaligned instructions. Our melding algorithm does not depend on the sequence alignment algorithm used for instruction alignment computation. We use Smith-Waterman algorithm [19] to compute the instruction alignment because prior work [5] has shown its effectiveness. Figure 3a shows the instruction alignment computed for two

basic blocks A and B . Aligned instructions are shown in green and instructions aligned with a gap are in red.

D. DARM Code Generation

Algorithm 1: DARM Algorithm

```

Input: SPMD function  $F$ 
Output: Melded SPMD function  $F_{out}$ 
do
  changed  $\leftarrow$  false
  for  $BB$  in  $F$  do
     $R, C \leftarrow$  GetRegionFor(BB)
    if IsMeldableDivergent( $R$ ) then
      SimplifyRegion( $R$ )
       $A \leftarrow$  ComputeSubgraphAlignment( $R$ )
      for  $(S_T, S_F, profit)$  in  $A$  do
        if  $profit \geq threshold$  then
          Meld( $S_T, S_F, C$ )
          changed  $\leftarrow$  true
        end
      end
    end
    if changed then
      SimplifyFunction( $F$ )
      RecomputeControlFlowAnalyses( $F$ )
      break
    end
  end
while changed;
```

Algorithm 2: SESE Subgraph melding Algorithm

```

Input: SESE subgraphs  $S_T, S_F$ , Condition  $C$ 
Output: Melded SESE subgraph  $S_{out}$ 
List blockPairs  $\leftarrow$  Linearize( $S_T, S_F$ )
List A  $\leftarrow$  empty
for  $(B_T, B_F)$  in blockPairs do
  List instrPairs  $\leftarrow$  ComputeInstrAlignment( $B_T, B_F$ )
  A.append(instrPairs)
end
PreProcess( $S_T, S_F$ )
Map operandMap  $\leftarrow$  empty
for  $P$  in A do
   $I_{melded} \leftarrow$  Clone( $P$ )
  Update(operandMap,  $I_{melded}$ ,  $P$ )
end
for  $P$  in A do
  SetOperands( $P$ , operandMap, C)
end
RunUnpredication()
RunPostOptimizations()
```

DARM’s control-flow melding procedure is shown in algorithm 1. This algorithm takes in a SPMD function F and iterates over all basic blocks in F to check if the basic block is an entry to a meldable divergent region (R) according to

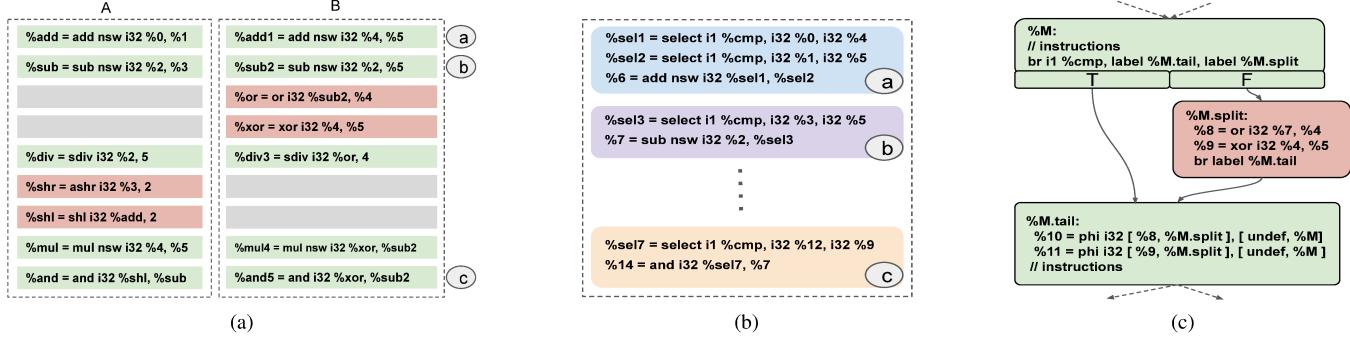


Fig. 3. (a) Instruction alignment result for two basic blocks A and B , (b) Code generated by DARM for aligned instructions ④, ⑤ and ⑥ in Figure 3a, (c) Unpredication applied to the unaligned instructions of basic block B in figure 3a

the conditions in Definition 5. We use *Simplify* to convert all subregions inside R in to simple regions.

We compute the optimal subgraph alignment for the two sequences of subgraphs in the true and false paths of R . We meld each subgraph pair in the alignment if the melding profitability is greater than some threshold. Subgraph melding changes the control-flow of F . Therefore we first simplify the control-flow (using LLVM’s *simplifycfg*) and then recompute the control-flow analyses (e.g. dominator, post-dominator and region tree) required for the melding pass. We apply the melding procedure on F again until no profitable melds can be performed.

Algorithm 2 shows the procedure for melding two subgraphs S_T and S_F . C is the branching condition of the meldable divergent region containing S_T and S_F . First the two subgraphs are linearized in pre-order to form a list of corresponding basic block pairs. Processing the basic blocks in pre-order ensures that dominating definitions are melded before their uses. For each basic block pair in this list we compute an optimal alignment of instructions. Each pair in the alignment falls into two categories, I - I and I - G . I - I is a proper alignment with two instructions and I - G is an instruction aligned with a gap. Our alignment makes sure that in a match the two instructions are always meldable into one instruction (e.g. a *load* is not allowed to align with a *store*). First we traverse the alignment pair list and clone the aligned instructions. For I - I pairs, we clone a single instruction because they can be melded. During cloning, we also update the *operandMap*, which maintains a mapping between aligned and melded LLVM values. We perform a second pass over the instruction alignment to set the operands of cloned instructions (*SetOperands*). Assume we are processing an I - I pair with instructions I_T, I_F and cloned instruction is I_{melded} . For each operand of I_{melded} , the corresponding operands from I_T and I_F are looked up in *operandMap* because an operand might be an already melded instruction. If the resultant two operands from I_T and I_F are the same, we just use that value as the operand. If they are different, we generate a *select* instruction to pick the correct operand conditioned by C . For an I - G pair, operands are first looked up in *operandMap* and the result is copied to I_{melded} . Consider the instruction alignment in figure 3a. Figure 3b

shows the generated code for aligned instruction pairs ④, ⑤ and ⑥. In case ④, two select instructions are needed because both operands maps to different values (%0, %4 and %1, %5). In case ⑤, the first operand is the same (%2) for both instructions, therefore only one select is needed. In case ⑥, both first and second operands are different for the two instructions. However the second operands map to same melded instruction %7, so only one select is needed. Note that %cmp is the branching condition for the divergent region, and we use that for selecting the operands.

Melding Branch Instructions of Exit Blocks: Setting operands for branch instructions in subgraph exit blocks is slightly different than that for other instructions. Let B_T^E, B_F^E be the exit blocks of S_T and S_F . Successors B_T^E, B_F^E can contain ϕ nodes. Therefore we need to ensure that successors of B_T^E and B_F^E can distinguish values produced in true path or false path. To solve this we move the branch conditions of B_T^E and B_F^E in to newly created blocks B'_T and B'_F . Now we can conditionally branch to B'_T and B'_F depending on C . For example, in Figure 4c basic blocks %M and %N are created when melding the exit branches of %X1 and %X2 in figure 4b. Any ϕ node in %G (figure 4c) can distinguish the values produced in true or false path using %M and %N.

Melding ϕ Nodes : In LLVM SSA form ϕ nodes are always placed at the beginning of a basic block. Even if the instruction alignment result contains two aligned ϕ nodes we can not meld them into a single ϕ node because *select* instructions can not be inserted before them. Therefore we copy all ϕ nodes into the melded basic block and set the operands for them using the *operandMap*. This can introduce redundant ϕ nodes which we remove during post-processing.

E. Unpredication

In our code generation process, unaligned instructions are inserted to the same melded basic block regardless of whether they are from true or false paths (i.e. fully predicated). This can introduce overhead due to several reasons. If the branching conditions C is biased towards the true or false path, it can result in redundant instruction execution. Also full predication of unaligned store instructions require adding extra loads to makes sure correct value is written back to

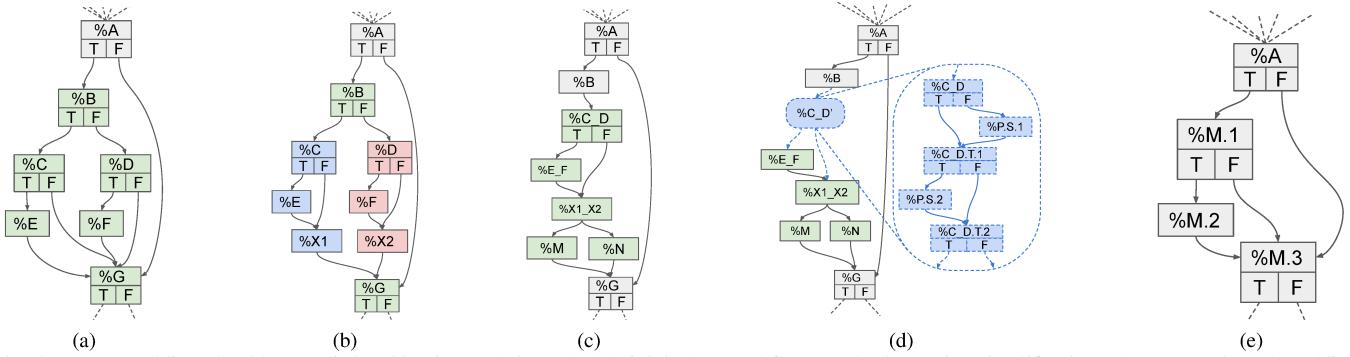


Fig. 4. DARM melding algorithm applied to bitonic sort (Figure 1) (a) Original control-flow graph, (b) Region simplification, (c) DARM subgraph melding, (d) Unpredication, (e) Final optimized control-flow graph

the memory. *Unpredication* splits the melded basic blocks at gap boundaries and moves the unaligned instructions into new blocks. Figure 3c shows unpredication applied to the unaligned instructions of basic block B in Figure 3a. The original basic block is split to two parts ($\%M$ and $\%M.tail$) and unaligned instructions ($\%8$ and $\%9$) are moved to a new basic block, $\%M.split$. ϕ nodes ($\%10$ and $\%11$) are added to $\%M.tail$ to ensure unaligned instructions dominate their uses. $\%8$ and $\%9$ are never executed in the true path, therefore ϕ nodes' incoming values from block $\%M$ are undefined (*LLVM undef*). Note that in region replication (Section IV-C) we apply unpredication only to the melded basic blocks. Store instructions outside the melded blocks are fully predicated by inserting extra loads.

F. Pre and Post Processing Steps

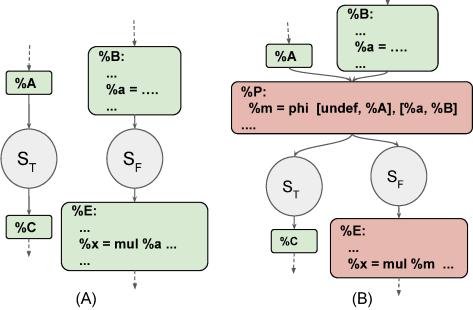


Fig. 5. DARM pre-processing example

In SSA form, any definition must dominate all its users. However DARM's subgraph melding can break this property. Consider the two meldable subgraphs S_T , S_F in figure 5 ①. Definition $\%a$ dominates its use $\%x$ before the melding. However if S_T and S_F are melded naively then $\%a$ will no longer dominate $\%x$. To fix this we add a new basic block $\%P$ with a ϕ node $\%m$. All uses of $\%a$ are replaced with $\%m$ (Figure 5 ②). Notice that value $\%m$ is never meant to be used in the true path execution. Therefore it is undefined in true path (*undef*). We apply this preprocessing step before the melding (*PreProcess* in Algorithm 2).

Subgraph melding can introduce branches with identical successors, ϕ nodes with identical operands and redundant ϕ nodes. *RunPostOptimizations* in Algorithm 2 removes these redundancies.

G. Putting All Together

Figure 4 shows how each stage of the pipeline of subgraph-melding transforms the CFG of bitonicSort kernel. The original CFG is shown in Figure 4a. Region $(\%B, \%G)$ is a meldable divergent region. Figure 4b shows the CFG after region simplification. Subgraphs $(\%C, \%X1)$ and $(\%D, \%X2)$ are profitable to meld according to our analysis. Figure 4c shows the CFG after subgraph-melding. The result after applying unpredication is shown in Figure 4d. Notice that the unpredication splits the basic block $\%C_D$ (in Figure 4c) into 5 basic blocks (zoomed in blue-dashed blocks in Figure 4d). Basic blocks $\%P.S.1$ and $\%P.S.2$ are the unaligned groups of instructions and they are executed conditionally. Figure 4e shows the final optimized CFG after applying post optimizations. Note that *ROCm HIPCC* compiler applied *if-conversion* aggressively. Therefore the effect of unpredication step is nullified in this case.

Figure 4 only shows how DARM transformation changes the CFG of our running example. It does not show how the instructions inside these transformed basic blocks are generated. We use Figure 6 to explain the generation of melded instructions for the running example. Figure 6a shows the LLVM-IR of the meldable divergent region $(\%B, \%G)$ in Figure 4b) in our running example. During DARM code generation, basic blocks in subgraphs $(\%C, \%X1)$ and $(\%D, \%X2)$ are linearized to compute the instruction alignment. Notice that $[\%C, \%D], [\%E, \%F], [\%X1, \%X2]$ are the corresponding basic block pairs. In this example all instructions perfectly align with each other except for the compare instructions ($\%34$ and $\%31$) in basic blocks $\%D$ and $\%C$. These compare instructions can not be aligned because their comparison kind is different (*greater than* vs *less than*). Figure 6b shows the LLVM-IR after applying subgraph melding and unpredication (similar to Figure 4d). Note that because instructions $\%34$ and $\%31$ are unaligned, unpredication step introduced basic blocks $\%P.S.1$ and $\%P.S.2$ to execute them conditionally based on the divergent condition $\%16$. Extra ϕ instructions $\%phi.1$ and $\%phi.2$ are inserted to ensure def-use chains are not broken during the unpredication step. Out of the all aligned instructions only the branch instructions at the end of basic blocks $\%C$ and $\%D$ require select instructions during instruction-melding. For example the store instructions in basic

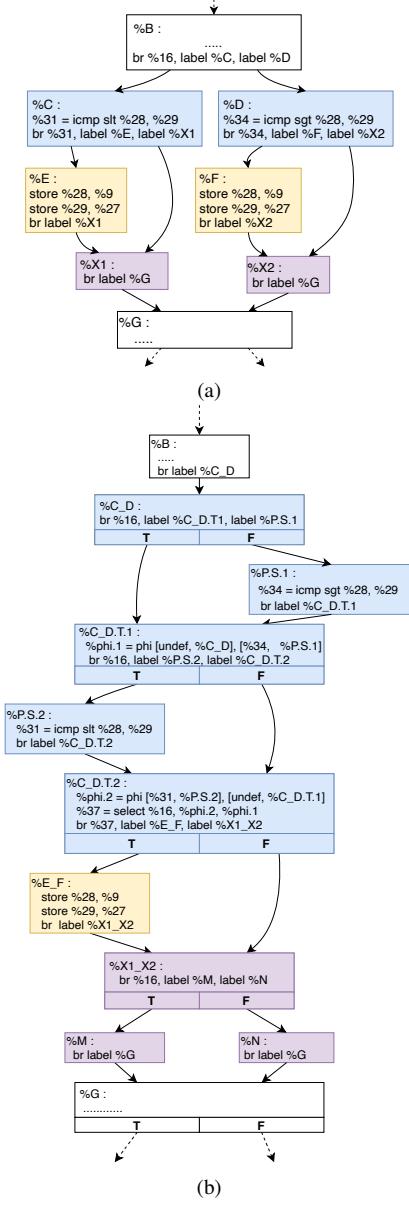


Fig. 6. LLVM-IR before and after applying DARM transformation to our running example (a) meldable divergent region (b) instruction alignment (b) LLVM-IR generated after subgraph melding and unpredication

blocks $\%E$, $\%F$ use matching operands, therefore can be melded without adding selects. On the other hand, conditional branch instructions uses values $\%34$ and $\%31$ and select instruction $\%37$ is inserted (Figure 6b) to pick the branching condition *conditionally*. Note that the values $\%34$ and $\%31$ will flow to their users via the ϕ nodes $\%phi.1$ and $\%phi.2$ respectively. Therefore the select instruction (*i.e.* $\%37$) uses these ϕ nodes as its operands.

V. IMPLEMENTATION

We implemented the DARM algorithm described in Section IV as an LLVM-IR analysis and transformation pass on top of the *ROCM HIPCC*³ GPU compiler [10]. Both the analysis and transformation are function passes that operate

³LLVM version 12.0.0, ROCm version 4.2.0

on GPGPU functions. The analysis pass first detects meldable divergent regions using LLVM's divergence analysis. Then it finds all the profitable subgraph pairs that can be melded. We use a default melding profitability threshold of 0.2 (algorithm 1). We also provide a sensitivity analysis on this threshold in Section VI-E. We use modified version of LLVM cost model [22] to obtain instruction latencies for melding profitability and instruction alignment computations. The transformation uses the output of analysis to perform DARM's code generation procedure (Section IV-D). The transformation pass also performs the unpredication, pre- and post-processing steps described in Sections IV-E and IV-F. LLVM pass is implemented in ~ 2500 lines of C++ code. In order to produce the program binary with our pass, we had to include our pass in the *ROCM HIPCC* compilation pipeline. Most GPGPU compilers (*e.g.* CUDA nvcc, ROCm HIPCC) use *separate compilation* for GPU device and CPU host codes. Final executable contains the device binary embedded in the host binary. In the modified workflow, we first compile the device code into LLVM-IR and run DARM on top of that to produce a transformed IR module. Our pass runs only on device functions and avoids any modifications to host code. After that, we use the LLVM static compiler (*llc*) [23] to generate an object file for the transformed device code. The rest of the compilation flow is as same as the one without any modification.

VI. EVALUATION

A. Evaluation Setup and Benchmarks

We evaluate the performance of DARM on a machine with a *AMD Radeon Pro Vega 20* GPU. This GPU has 16 GBs of global memory, 64 kB of shared memory (*i.e.* Local Data Share (LDS)) and 1700 MHz of max clock frequency. The machine consists of *AMD Ryzen Threadripper 3990X* 64-Core Processor with 2900 MHz max clock frequency.

We use two different sets of benchmarks. First, to assess the generality of DARM, we create several synthetic programs that exhibit control divergence of varying complexity. While many real-world programs are hand-optimized to eliminate divergence, these synthetic programs qualitatively demonstrate the generality of DARM over prior automated divergence-reduction techniques, and show that DARM can automate the control-flow melding that would be otherwise done by hand. For detailed description of the evaluation on synthetic benchmarks, please refer to the extended version of our paper [24].

Real-world Benchmarks We show DARM's effectiveness on real-world programs. We consider 7 benchmarks written in *HIP* [7]. These benchmarks were taken from well-known highly hand-optimized GPU benchmark suites or optimized reference implementations of papers. We selected these benchmarks because they contain divergent if-then-else regions that present melding opportunities for DARM. We do not consider benchmarks that do not present any melding opportunities for DARM because they are not modified by DARM in any way.

Bitonic Sort (BIT) Our running example is bitonic sort [14]. In this kernel, each thread block takes in a bucket

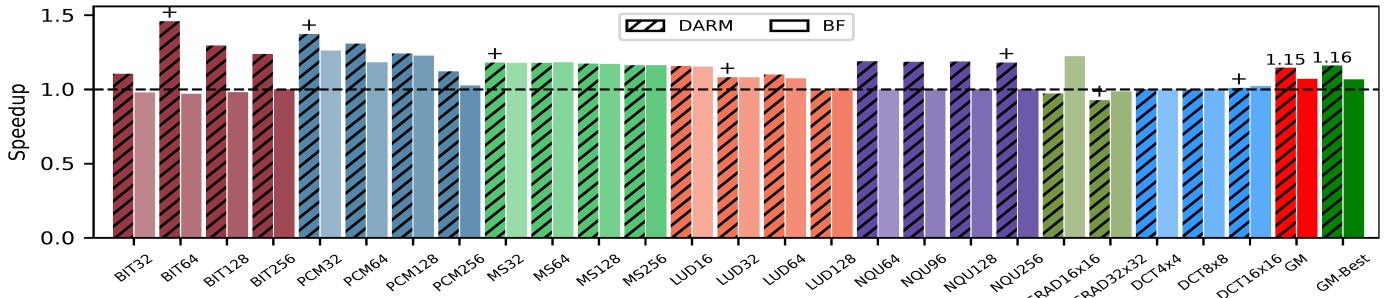


Fig. 7. Real-world Benchmark Performance. + marks block size with best baseline runtime. GM is geo-mean of DARM’s speedup on all benchmarks; GM-Best is DARM’s speedup on + configurations.

and performs parallel sort. We used an input of 2^{26} elements and varied the bucket (*i.e.* block) size.

Partition and Concurrent Merge (PCM) PCM is a parallel sorting algorithm based on Batcher’s odd-even merge sort [25]. PCM performs odd-even merging of *buckets* of sorted elements at every position of the array leading to loops with nested data-dependent branches. We used an array of 2^{28} elements with different number of buckets.

Mergesort (MS) A parallel bottom-up merge sort implementation. The kernel has data-dependent control-flow divergence in the merging step. We used an input array with 2^{20} elements.

LU-Decomposition (LUD) LUD implementation from the Rodinia benchmark suite [26]. We focus our evaluation on the *lud_perimeter* kernel in this benchmark. *lud_perimeter* contains multiple divergent branches that depend on thread ID and block size. We use a randomly generated matrix of size 16384×16384 as the input.

N-Queens (NQU) N-Queens solver uses backtracking to find all different ways of placing N queens on a NxN chessboard without attacking each other. We have used the kernel from the GPGPU-sim benchmark suite [27] with N is 15.

Speckle Reducing Anisotropic Diffusion (SRAD) SRAD is diffusion based noise removal method for imaging applications from Rodinia benchmark suite [26]. We have used an image of size 4096×4906 as input.

DCT Quantization (DCT) An in-place quantization of a discrete cosine transformation (DCT) plane [28]. The quantization process is different for positive and negative values resulting in data-dependent divergence. We use a randomly generated DCT plane of size $2^{15} \times 2^{15}$ as input.

Baseline and Branch Fusion Our baseline implementations of these kernels have been hand-optimized (except, obviously, for optimizations that manually remove control divergence by applying DARM-like transformations). This optimization includes using shared memory when needed to improve performance. The baseline implementations were compiled with -O3. Branch fusion [5] was implemented in the Ocelot [29] open-source CUDA compiler that is no longer maintained and does not support AMD GPUs. We implemented branch fusion by modifying DARM to apply melding for diamond-shaped control-flow (*if-then-else*). We use this for comparison against branch fusion. Branch fusion cannot fully handle the

control-flow of BIT, PCM, and NQU. Loop unrolling enables successful branch fusion in LUD.

Block Size Each of these kernels has a tunable *block size*—essentially, a tile size that controls the granularity of work in the inner loops. Because the correct block size can be dependent on many parameters (though for a given input and GPU configuration, one is likely the best), our evaluation treats block size as exogenous to the evaluation, and hence considers behavior at different block sizes for each kernel. In other words, our evaluation asks: if a programmer has a kernel with a given block size, what will happen if DARM is applied?

Note that of these kernels, only LUD exhibit divergence that depends on block size. This means that all the other benchmarks will experience divergence regardless of block size. LUD’s divergence, on the other hand, is block size dependent. For some block sizes, the kernel will be divergent, while for others, it will be convergent.

B. Performance

Figure 7 shows the speedups for real benchmarks DARM always improves the performance ($1.15 \times$ geo-mean speedup over all benchmarks and $1.16 \times$ geo-mean speedup over the best baseline variants) except for SRAD (see below). The highest relative improvement in performance can be seen in BIT and PCM for all block sizes. This is because both these benchmarks are divergent regardless of the block size and they have complex control-flow regions with shared memory instructions. DARM successfully melds these regions and reduces divergence significantly. Branch fusion improves performance in PCM by melding *if-then-else* blocks. In LUD, the divergence is block size dependent, and the kernel is divergent only at block sizes 16, 32 and 64, where we see a visible performance improvement introduced by DARM. NQU contains a time-consuming loop with divergent *if-then-elseif-then* section. DARM applies *region replication* to remove divergence, achieving superior performance. SRAD kernel has both block size-dependent and data-dependent divergent regions (say R_B and R_D respectively). Both R_B and R_D consists of *if-then-else-if-then-else* chains. R_B contains no shared memory instructions and melding does not improve performance (for both DARM and branch fusion). However R_D contains a 3-way divergent branch with shared memory instructions and the divergence is biased *i.e.* execution only takes 2 of the 3 ways. In this case branch fusion has better

performance at block size 16, because blocks that get melded happen to be on the divergent paths. However DARM has more melding options than branch fusion, and it melds all 3 paths adding extra overhead. At block size 32, the extra overhead introduced by melding R_B becomes significant and both DARM and branch fusion exhibit a performance drop. Performance drop for DARM can be avoided by prioritizing the melding order (*i.e.* apply melding to divergent regions with most profitable subgraphs first). However, prioritizing melding order is not considered in this paper.

In most cases (except SRAD), the block size for best performing baseline is also the one that gives the best absolute performance for DARM. Interestingly, for 4/7 benchmarks (BIT, PCM, MS, and DCT), not only does this best baseline block size produce the best absolute DARM performance, it also produces the best *speedup* relative to the baseline: the block size that makes the baseline perform the best, actually exposes more optimization opportunities to DARM.

We use *rocprof* [30] to collect ALU utilization and memory instruction counters to reason about performance. We focus on the block sizes for each benchmark where DARM has highest improvement over the baseline.

C. ALU Utilization

DARM’s melding transformation enables the ALU instructions in divergent paths to be issued in the same cycle. This effectively improves the SIMD resource utilization. Figure 8 shows the ALU utilization (%). As expected DARM improves the ALU utilization significantly for most benchmarks. In BIT, divergent paths does not have common comparison operators ($>$ and $<$ comparisons in lines 9 and 13 in Figure 1). Even though DARM unpredicts these instructions, later optimization passes decide to fully-predicate them resulting in lower ALU utilization.

D. Melding of Memory Instructions

Figure 9 shows the normalized number of global and shared memory (*i.e.* local data share) instructions issued after apply-

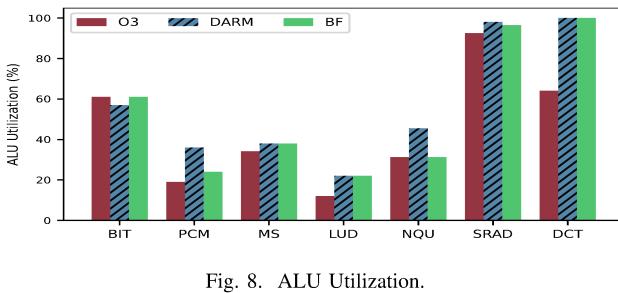


Fig. 8. ALU Utilization.

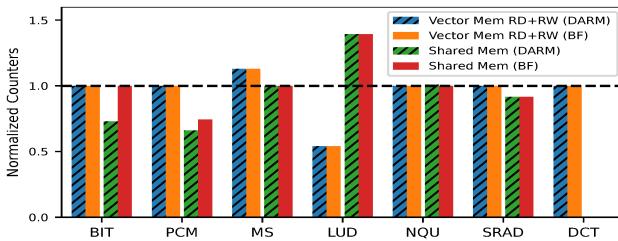


Fig. 9. Normalized Memory Instruction Counters.

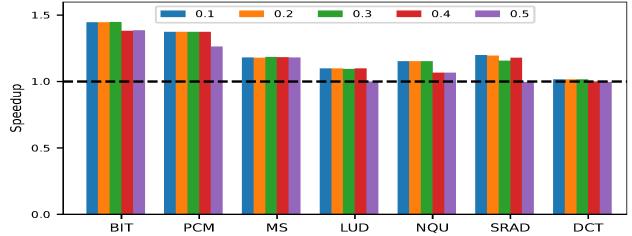


Fig. 10. Variation of melding profitability thresholds.

ing DARM. In LUD, there are many common shared memory instructions in divergent paths. However these instructions do not have different memory alignments, therefore cannot be melded into a single instruction. Unpredicted shared memory instructions are predicted by other optimization passes in LLVM resulting in higher instruction count. Melding reduces the global memory instruction count in LUD. DCT does not have any memory instructions in the divergent region and does not use shared memory. In BIT and PCM, the melded regions contain a lot of shared memory instructions. Therefore the reduction in shared memory instructions is significant and correlate with the performance gain. We find that melding shared memory instructions is more beneficial than melding ALU instructions because shared memory instructions have higher latency than most ALU instructions, though lower latency than global memory instructions. Therefore there is 2 \times improvement in cycles spent if two divergent shared memory instructions are issued in the same cycle. In contrast, melding global memory instructions does not always improve performance. This is because the data requested by divergent memory instructions might be on different cache lines and these requests are serialized by the memory controller even if they are issued in the same cycle.

E. Melding Profitability Threshold

Figure 10 shows the performance of DARM for different melding profitability thresholds on the real-world benchmarks considering DARM’s best performing block sizes. For all benchmarks, we observe that DARM’s speedup reduces as we increase the threshold due to lost opportunities. When we reduce the threshold, increment in the improvement of the performance of DARM becomes insignificant (after 0.2). But we cannot reduce it to zero because every possible pair would be melded and the subsequent CFG simplification passes would unpredicate them. As a result, DARM may become non-convergent.

F. Compile Time

Benchmark	AVERAGE COMPILE TIME (S)		
	O3	DARM	Normalized
BIT	0.4804	0.5018	1.0444
PCM	0.5690	0.5942	1.0443
MS	0.8037	0.8064	1.0035
LUD	0.5993	0.6294	1.0502
NQU	0.4687	0.4738	1.0109
SRAD	0.4999	0.5121	1.0244
DCT	0.4398	0.4439	1.0093

Table II shows the device code compilation times for the baseline and DARM. We omit the time for compiling host

code and linking because it is constant for both the baseline and DARM. Since we perform the analysis and the instruction alignment – the most costly parts – at the basic block level rather than performing at a higher level (*i.e.* function or region level), we incur negligible compilation overhead. Compilation time overhead introduced by DARM is a small fraction of total compilation time (including host code) for all cases.

DARM’s compile time depends on the size of basic blocks that get melded and the structure of the program since it determines different types of melding opportunities. A slight overhead in compilation time of LUD is caused by sequence alignment overhead on large basic blocks (created by loop unrolling). PCM and BIT have divergent regions inside an unrolled loop, therefore DARM’s meldable subgraph detection incurs overhead. Only BIT and PCM has opportunities for region to region melding, and only PCM, NQU, and SRAD have opportunities for basic block to region melding. Presence of basic block to region melding opportunity results in *region replication*.

VII. RELATED WORK

Impact of control-flow divergence has extensively studied in different contexts [31]–[34]. Reducing control-flow divergence requires finding the source of divergence in a program. Coutinho *et al.* constructed a divergence analysis to statically identify variables with the same value for every SIMD unit and used this analysis to drive *Branch Fusion* [5]. A divergence analysis of similar fashion based on data and sync dependences has been integrated to the LLVM framework [12]. Recently, Rosemann *et al.* has presented a precise divergence analysis based on abstract interpretation for reducible CFGs [13]. Using a precise divergence analysis improves the opportunities of melding for DARM.

Tail Merging is a standard, but restrictive, compiler optimization used to reduce the code size by merging identical sequences of instructions. Chen *et al.* used generalized tail merging to compact matching *Single-Entry-Multiple-Exit* regions [4]. Recently, Rocha *et al.* has presented *Function Merging*, an advanced sequence-alignment based technique for code size reduction [20], [21]. Even though parts of DARM has some similarities with function merging, it does not tackle divergence.

In addition to branch fusion, Anantpur and Govindarajan proposed to structure the unstructured CFGs and then linearize it with predication [35]. More recently, Fukuhara and Takimoto proposed *Speculative Sparse Code Motion* to reduce divergence in GPU programs [36], which preserves the CFG and it is orthogonal to DARM. *Collaborative Context Collection* copies registers of divergent warps to shared memory and restores them when those warps become non-divergent [37]. *Iteration Delaying* is a complementary compiler optimization to DARM that delays divergent loop iterations [38] and can be applied following DARM. Recently, Damani *et al.* has presented a speculative reconvergence technique for GPUs similar to iteration delaying [39]. *Common Subexpression Convergence (CSC)* [40] works similar to branch fusion but

uses *branch flattening* (*i.e.* predication) to handle complex control-flow. In contrast, DARM does not require predication to meld complex control-flow, thus more general than CSC.

Architectural techniques such as Thread Block Compaction [41] and Dynamic Warp Formation [1] involve repacking threads into non-divergent warps. Variable Warp Sizing [42] and Dynamic Warp Subdivision [43] depend on smaller warps to schedule divergent thread groups in parallel. Independent Thread Scheduling helps to hide the latency in divergent paths by allowing to switch between divergent threads inside a warp [3], [44].

VIII. DISCUSSION AND FUTURE WORK

Most of the GPGPU benchmarks are heavily hand optimized by expert developers and this often include DARM like transformations to remove control-flow divergence [5]. We evaluate DARM on limited set of real-world benchmarks mainly because of this reason. However we also emphasize that doing DARM-like transformations by hand is time-consuming and error-prone. For example, it took us several hours to manually apply control-flow melding to LUD kernel. Therefore, offloading this to the compiler can save a lot of developer effort.

The benefits of DARM is not limited to reducing control-flow divergence in GPGPU programs. DARM can be used to reduce control-flow divergence in any hardware backends and programming models that employ SIMD execution (*e.g.* intel/AMD processors with ISPC [45]). DARM can be used to reduce branches in a program. This property can be exploited to accelerate software testing techniques such as symbolic execution [46]. DARM factor out common code segments within *if-the-else* regions of a program. Therefore it can be used as an intra-function code size reduction optimization as well. Aforementioned applications of DARM suggest that it is useful as a general compiler optimization technique. We plan to explore some of these applications in our future work.

IX. CONCLUSION

Divergent control-flow in GPGPU programs causes performance degradation due to serialization. We presented DARM, a new compiler analysis and transformation framework for GPGPU programs implemented on LLVM, that can detect and meld similar control-flow regions in divergent paths to reduce divergence in control-flow. DARM generalizes and subsumes prior efforts at reducing divergence such as tail merging and branch fusion. We showed that DARM improves performance by improving ALU utilization and promoting coalesced shared memory accesses across several real-world benchmarks.

ACKNOWLEDGMENTS

This work was supported in part by National Science Foundation awards CCF-1919197 and CCF-1908504. We would like to thank Tim Rogers, Rodrigo Rocha and anonymous reviewers for their help during various stages of this work.

REFERENCES

[1] W. W. L. Fung, I. Sham, G. Yuan, and T. M. Aamodt, "Dynamic warp formation and scheduling for efficient gpu control flow," in *40th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO 2007)*, 2007, pp. 407–420.

[2] W. W. L. Fung and T. M. Aamodt, "Thread block compaction for efficient simt control flow," in *2011 IEEE 17th International Symposium on High Performance Computer Architecture*, 2011, pp. 25–36.

[3] M. Rhu and M. Erez, "The dual-path execution model for efficient gpu control flow," in *2013 IEEE 19th International Symposium on High Performance Computer Architecture (HPCA)*, 2013, pp. 591–602.

[4] W.-K. Chen, B. Li, and R. Gupta, "Code compaction of matching single-entry multiple-exit regions," in *Proceedings of the 10th International Conference on Static Analysis*, ser. SAS'03. Berlin, Heidelberg: Springer-Verlag, 2003, p. 401–417.

[5] B. Coutinho, D. Sampaio, F. M. Q. Pereira, and W. Meira Jr., "Divergence analysis and optimizations," in *2011 International Conference on Parallel Architectures and Compilation Techniques*, 2011, pp. 320–329.

[6] C. Lattner and V. Adve, "Llvm: a compilation framework for lifelong program analysis transformation," in *International Symposium on Code Generation and Optimization, 2004. CGO 2004.*, 2004, pp. 75–86.

[7] "HIP Programming Guide v4.1," [Accessed 17-Dec-2021]. [Online]. Available: <https://rocmdocs.amd.com/en/latest/>

[8] "CUDA C++ Programming Guide," [Accessed 17-Dec-2021]. [Online]. Available: <https://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html>

[9] "NVCC :: CUDA Toolkit Documentation," [Accessed 17-Dec-2021]. [Online]. Available: <https://docs.nvidia.com/cuda/cuda-compiler-driver-nvcc/index.html>

[10] "ROCM Compiler SDK," [Accessed 17-Dec-2021]. [Online]. Available: https://rocmdocs.amd.com/en/latest/ROCM_Compiler_SDK/ROCM-Compiler-SDK.html

[11] R. Cytron, J. Ferrante, B. K. Rosen, M. N. Wegman, and F. K. Zadeck, "Efficiently computing static single assignment form and the control dependence graph," *ACM Trans. Program. Lang. Syst.*, vol. 13, no. 4, p. 451–490, Oct. 1991. [Online]. Available: <https://doi.org/10.1145/115372.115320>

[12] R. Karrenberg and S. Hack, "Improving performance of opencl on cpus," in *Compiler Construction*, M. O'Boyle, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 1–20.

[13] J. Rosemann, S. Moll, and S. Hack, "An abstract interpretation for spmvd divergence on reducible control flow graphs," *Proc. ACM Program. Lang.*, vol. 5, no. POPL, Jan. 2021. [Online]. Available: <https://doi.org/10.1145/3434312>

[14] K. E. Batcher, "Sorting networks and their applications," in *Proceedings of the April 30–May 2, 1968, spring joint computer conference (AFIPS '68 (Spring))*, 1968, p. 307–314.

[15] D. Cederman and P. Tsigas, "Gpu-quicksort: A practical quicksort algorithm for graphics processors," *ACM J. Exp. Algorithmics*, vol. 14, Jan. 2010. [Online]. Available: <https://doi.org/10.1145/1498698.1564500>

[16] "llvm::RegionBase Class Template Reference," [Accessed 17-Dec-2021]. [Online]. Available: https://llvm.org/doxygen/classllvm_1_1RegionBase.html

[17] R. Johnson, D. Pearson, and K. Pingali, "The program structure tree: Computing control regions in linear time," *SIGPLAN Not.*, vol. 29, no. 6, p. 171–185, Jun. 1994. [Online]. Available: <https://doi.org/10.1145/773473.178258>

[18] "Using cuda warp-level primitives," [Accessed 17-Dec-2021]. [Online]. Available: <https://developer.nvidia.com/blog/using-cuda-warp-level-primitives/>

[19] T. Smith and M. Waterman, "Identification of common molecular subsequences," *Journal of Molecular Biology*, vol. 147, no. 1, pp. 195–197, 1981. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0022283681900875>

[20] R. C. O. Rocha, P. Petoumenos, Z. Wang, M. Cole, and H. Leather, "Function merging by sequence alignment," in *2019 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*, 2019, pp. 149–163.

[21] R. C. O. Rocha, P. Petoumenos, Z. Wang, M. Cole, and H. Leather, "Effective function merging in the ssa form," in *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI 2020. New York, NY, USA: Association for Computing Machinery, 2020, p. 1–12. [Online]. Available: <https://doi.org/10.1145/3356559.3385603>

[22] "CostModel.cpp File Reference," [Accessed 17-Dec-2021]. [Online]. Available: https://llvm.org/doxygen/CostModel_8cpp.html

[23] "llc - LLVM static compiler," [Accessed 17-Dec-2021]. [Online]. Available: <https://llvm.org/docs/CommandGuide/llc.html>

[24] C. Saumya, K. Sundararajah, and M. Kulkarni, "Darm: Control-flow melding for simt thread divergence reduction - extended version," 2021.

[25] E. Herruzo, G. Ruiz, J. I. Benavides, and O. Plata, "A new parallel sorting algorithm based on odd-even mergesort," in *15th EUROMICRO International Conference on Parallel, Distributed and Network-Based Processing (PDP'07)*, 2007, pp. 18–22.

[26] S. Che, M. Boyer, J. Meng, D. Tarjan, J. W. Sheaffer, S. Lee, and K. Skadron, "Rodinia: A benchmark suite for heterogeneous computing," in *2009 IEEE International Symposium on Workload Characterization (IISWC)*, 2009, pp. 44–54.

[27] A. Bakhoda, G. L. Yuan, W. W. L. Fung, H. Wong, and T. M. Aamodt, "Analyzing cuda workloads using a detailed gpu simulator," in *2009 IEEE International Symposium on Performance Analysis of Systems and Software*, 2009, pp. 163–174.

[28] "CUDA Samples," [Accessed 17-Dec-2021]. [Online]. Available: <https://docs.nvidia.com/cuda/cuda-samples/>

[29] A. Kerr, G. Diamos, and S. Yalamanchili, "A characterization and analysis of ptx kernels," in *2009 IEEE International Symposium on Workload Characterization (IISWC)*, 2009, pp. 3–12.

[30] "ROCM-Developer-Tools / rocprofiler," [Accessed 17-Dec-2021]. [Online]. Available: <https://github.com/ROCM-Developer-Tools/rocprofiler>

[31] T. Schaub, S. Moll, R. Karrenberg, and S. Hack, "The impact of the simd width on control-flow and memory divergence," *ACM Trans. Archit. Code Optim.*, vol. 11, no. 4, Jan. 2015. [Online]. Available: <https://doi.org/10.1145/2687355>

[32] R. Karrenberg and S. Hack, "Whole Function Vectorization," in *International Symposium on Code Generation and Optimization*, ser. CGO, 2011. [Online]. Available: http://www.cdl.uni-saarland.de/papers/karrenberg_wfv.pdf

[33] S. Moll and S. Hack, "Partial Control-flow Linearization," in *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI 2018. New York, NY, USA: ACM, 2018, pp. 543–556. [Online]. Available: <http://doi.acm.org/10.1145/3192366.3192413>

[34] T. Lloyd, K. Ali, and J. N. Amaral, "Gpucheck: Detecting cuda thread divergence with static analysis," Department of Computer Science, University of Alberta, Tech. Rep., 2019. [Online]. Available: <https://era.library.ualberta.ca/items/7ab2b28d-b111-448f-8273-2ff219132908>

[35] J. Anantpur and G. R., "Taming control divergence in gpus through control flow linearization," in *Compiler Construction*, A. Cohen, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 133–153.

[36] J. Fukuhara and M. Takimoto, "Branch divergence reduction based on code motion," *Journal of Information Processing*, vol. 28, pp. 302–309, 2020.

[37] F. Khorasani, R. Gupta, and L. N. Bhuyan, "Efficient warp execution in presence of divergence with collaborative context collection," in *Proceedings of the 48th International Symposium on Microarchitecture*, ser. MICRO-48. New York, NY, USA: Association for Computing Machinery, 2015, p. 204–215. [Online]. Available: <https://doi.org/10.1145/2830772.2830796>

[38] T. D. Han and T. S. Abdelrahman, "Reducing branch divergence in gpu programs," in *Proceedings of the Fourth Workshop on General Purpose Processing on Graphics Processing Units*, ser. GPGPU-4. New York, NY, USA: Association for Computing Machinery, 2011. [Online]. Available: <https://doi.org/10.1145/1964179.1964184>

[39] S. Damani, D. R. Johnson, M. Stephenson, S. W. Keckler, E. Yan, M. McKeown, and O. Giroux, "Speculative reconvergence for improved simt efficiency," in *Proceedings of the 18th ACM/IEEE International Symposium on Code Generation and Optimization*, ser. CGO 2020. New York, NY, USA: Association for Computing Machinery, 2020, p. 121–132. [Online]. Available: <https://doi.org/10.1145/3368826.3377911>

[40] S. Damani and V. Sarkar, "Common subexpression convergence: A new code optimization for simt processors," in *Languages and Compilers for Parallel Computing*, S. Pande and V. Sarkar, Eds. Cham: Springer International Publishing, 2021, pp. 64–73.

[41] W. W. L. Fung and T. M. Aamodt, "Thread block compaction for efficient simt control flow," in *2011 IEEE 17th International Symposium on High Performance Computer Architecture*, 2011, pp. 25–36.

[42] T. G. Rogers, D. R. Johnson, M. O'Connor, and S. W. Keckler, "A variable warp size architecture," in *Proceedings of the 42nd Annual International Symposium on Computer Architecture*, ser. ISCA '15. New York, NY, USA: Association for Computing Machinery, 2015, p. 489–501. [Online]. Available: <https://doi.org/10.1145/2749469.2750410>

[43] J. Meng, D. Tarjan, and K. Skadron, "Dynamic warp subdivision for integrated branch and memory divergence tolerance," *SIGARCH Comput. Archit. News*, vol. 38, no. 3, p. 235–246, Jun. 2010. [Online]. Available: <https://doi.org/10.1145/1816038.1815992>

[44] A. ElTantawy, J. W. Ma, M. O'Connor, and T. M. Aamodt, "A scalable multi-path microarchitecture for efficient gpu control flow," in *2014 IEEE 20th International Symposium on High Performance Computer Architecture (HPCA)*, 2014, pp. 248–259.

[45] M. Pharr and W. R. Mark, "ispcc: A spmd compiler for high-performance cpu programming," in *2012 Innovative Parallel Computing (InPar)*, 2012, pp. 1–13.

[46] C. Cadar, "Targeted program transformations for symbolic execution," in *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering*, ser. ESEC/FSE 2015. New York, NY, USA: Association for Computing Machinery, 2015, p. 906–909. [Online]. Available: <https://doi.org/10.1145/2786805.2803205>